

WILLIAM WEBBER

29 Canning Street, North Melbourne, Victoria 3051, Australia
william@williamwebber.com m: (+61|0)418 532 084

ACADEMIC QUALIFICATIONS

Ph. D. University of Melbourne, Australia, 2010. Thesis title: “Measurement in Information Retrieval Evaluation”. Supervisors: Alistair Moffat and Justin Zobel.

M. Eng. Sc. (Comp. Sc. & Soft. Eng.) (by research) University of Melbourne, Australia, 2006. Thesis title: “Design and Implementation of a Pipelined Distributed Information Retrieval Architecture”. Supervisors: Alistair Moffat and Justin Zobel.

Grad. Dip. Comp. Sc. RMIT University, Melbourne, Australia, 1997.

B. Arts (Honours) (First Class) University of Melbourne, Australia, 1992.

PROFESSIONAL EXPERIENCE

Private Consultant

August 2013 to present

William Webber Consulting

Consults with software vendors, service providers, system integrators, and end-users on development, process control, and statistical validation of e-discovery, information retrieval, and text analysis technologies.

Post-doctoral Research Associate

August 2011 to July 2013

College of Information Studies, University of Maryland, United States

Designed and validated statistical tests for e-discovery evaluation. Developed testing and evaluation software for information retrieval and text classification. Designed algorithms for graduate researchers, and supervised implementation. Investigated nature and causes of relevance assessor errors. Communicated results to academia and industry through peer-reviewed research publications, talks, webinars, blog, personal discussions, and collaborations. Taught graduate course in information technology. Supervised graduate research students in text classification statistical testing, document privileged detection, and information retrieval evaluation. Member of E-Discovery Lab and Computational Linguistics and Information Processing (CLIP) Lab. Research work performed as part of NSF-funded project.

Research Associate

February 2005 to July 2011

Department of Computer Science and Software Engineering, University of Melbourne, Australia

Designed statistical tests and evaluation metrics for the effectiveness of information retrieval and e-discovery systems. Led collaborative academic–industry team at the TREC Legal Track for two years. Designed and supervised development of online database of experimental results in information retrieval. Designed and implemented a high-performance cluster-based distribution information retrieval system. Communicated results through peer-reviewed research publications, talks, and blog. Taught graduate-level course in compression and information retrieval; undergraduate course in assembler and C programming. Supervised summer students in information retrieval software development; honors student in text classification research.

Research Assistant

January 2004 to February 2005

School of Computer Science and Information Technology, RMIT University, Melbourne, Australia

Designed and implemented a distributed information retrieval system. Supervised experiments for TREC participation.

Freelance Software Developer

July 2001 to January 2004

Various

Developed web applications in Java and PHP. Taught six-week industry course in C++ as sessional instructor at RMIT University.

- Senior Programmer** January 1999 to April 2001
PeopleWeb Communications, Melbourne, Australia
 Developed multi-function web portal software in Java, include chat, calendaring, and email applications.
- Tutor, Senior Tutor** March to December 1998
RMIT University, Melbourne, Australia
 Tutored postgraduate subjects in algorithms and data structures, and C programming. Managed other tutors as senior tutor.
- Research Programmer** July to December 1997
Multimedia Database Systems, RMIT University, Melbourne, Australia
 Performed development on the MG research full-text database system in C. Added threading, 64-bit support.
- Programmer** July to December 1996
Cincom Systems of Australia, Melbourne, Australia
 Developed enterprise inventory and customer relationship management software for leading international company in COBOL, 4GL, and REXX.

PUBLICATIONS

Books, chapters, and monographs

Douglas W. Oard and William Webber. Information retrieval for e-discovery. *Foundations and Trends in Information Retrieval*, 7(2–3):99–237, 2013.

Peer-reviewed journals

William Webber. Approximate recall confidence intervals. *ACM Transactions on Information Systems*, 31(1):2:1–33, 2013.

William Webber, Alistair Moffat, and Justin Zobel. A similarity measure for indefinite rankings. *ACM Transactions on Information Systems*, 28(4):20:1–38, November 2010a.

Alistair Moffat, William Webber, Justin Zobel, and R. Baeza-Yates. A pipelined architecture for distributed text query evaluation. *Information Retrieval*, 10(3):205–231, June 2007a.

Peer-reviewed conferences

Gaya K. Jayasinghe, William Webber, Mark Sanderson, Lasitha S. Dharmasena, and J. Shane Culpepper. Evaluating non-deterministic retrieval systems. In Peter Bruza, Charles L. A. Clarke, and Kalervo Järvelin, editors, *Proc. 37th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Gold Coast, Australia, July 2014a. To appear.

Gaya K. Jayasinghe, William Webber, Mark Sanderson, and J. Shane Culpepper. Extending test collection pools without manual runs. In Peter Bruza, Charles L. A. Clarke, and Kalervo Järvelin, editors, *Proc. 37th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Gold Coast, Australia, July 2014b. To appear.

Ning Gao, William Webber, and Douglas W. Oard. Reducing reliance on relevance judgments for system comparison by using expectation-maximization. In *Proc. 31st European Conference on IR Research*, pages 1–12, Amsterdam, the Netherlands, April 2014.

Mossaab Bagdouri, William Webber, David D. Lewis, and Douglas W. Oard. Towards minimizing the annotation cost of certified text classification. In *Proc. 22st ACM International Conference on Information and*

Knowledge Management, pages 989–998, San Francisco, CA, USA, October 2013.

William Webber and Jeremy Pickens. Assessor disagreement and text classifier accuracy. In Diane Kelly, Maarten de Rijke, and Tetsuya Sakai, editors, *Proc. 36th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Dublin, Ireland, July 2013. 929–932.

Praveen Chandar, William Webber, and Ben Carterette. Document features predicting assessor disagreement. In Diane Kelly, Maarten de Rijke, and Tetsuya Sakai, editors, *Proc. 36th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Dublin, Ireland, July 2013. 745–748.

William Webber, Mossaab Bagdouri, David D. Lewis, and Douglas W. Oard. Sequential testing in classifier evaluation yields biased estimates of effectiveness. In Diane Kelly, Maarten de Rijke, and Tetsuya Sakai, editors, *Proc. 36th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Dublin, Ireland, July 2013. 933–936.

Falk Scholer, Diane Kelly, Hanseul S. Lee, Wan-ching Wu, and William Webber. The effect of threshold priming and need for cognition on relevance assessment. In Diane Kelly, Maarten de Rijke, and Tetsuya Sakai, editors, *Proc. 36th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 623–632, Dublin, Ireland, July 2013.

William Webber, Praveen Chandar, and Ben Carterette. Alternative assessor disagreement and retrieval depth. In *Proc. 21st ACM International Conference on Information and Knowledge Management*, pages 125–134, Maui, Hawaii, USA, October 2012a.

William Webber, Bryan Toth, and Marjorie Desamito. Effect of written instructions on assessor agreement. In William Hersh, Jamie Callan, Yoelle Maarek, and Mark Sanderson, editors, *Proc. 35th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1053–1054, Portland, Oregon, USA, August 2012b.

William Webber, Douglas W. Oard, Falk Scholer, and Bruce Hedin. Assessor error in stratified evaluation. In *Proc. 19th ACM International Conference on Information and Knowledge Management*, pages 539–548, Toronto, Canada, October 2010b.

Timothy G. Armstrong, Alistair Moffat, William Webber, and Justin Zobel. Improvements that don't add up: Ad-hoc retrieval results since 1998. In David Cheung, Il-Yeol Song, Wesley Chu, Xiaohua Hu, Jimmy Lin, Jiexun Li, and Zhiyong Peng, editors, *Proc. 18th ACM International Conference on Information and Knowledge Management*, pages 601–610, Hong Kong, China, November 2009a.

William Webber and Laurence A. F. Park. Score adjustment for correction of pooling bias. In James Allan, Javed Aslam, Mark Sanderson, ChengXiang Zhai, and Justin Zobel, editors, *Proc. 32nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 444–451, Boston, Massachusetts, USA, July 2009.

Timothy G. Armstrong, Alistair Moffat, William Webber, and Justin Zobel. Has adhoc retrieval improved since 1994? In James Allan, Javed Aslam, Mark Sanderson, ChengXiang Zhai, and Justin Zobel, editors, *Proc. 32nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 692–693, Boston, Massachusetts, USA, July 2009b.

William Webber, Alistair Moffat, and Justin Zobel. Statistical power in retrieval experimentation. In James G. Shanahan, Sihem Amer-Yahia, Yi Zhang, Alek Kolcz, Abdur Chowdhury, and Diane Kelly, editors, *Proc. 17th ACM International Conference on Information and Knowledge Management*, pages 571–580, Napa, California, USA, October 2008a.

William Webber, Alistair Moffat, and Justin Zobel. Score standardization for inter-collection comparison of

retrieval systems. In Sung-Hyon Myaeng, Douglas W. Oard, Fabrizio Sebastiani, Tat-Seng Chua, and Mun-Kew Leong, editors, *Proc. 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 51–58, Singapore, Singapore, July 2008b.

William Webber, Alistair Moffat, Justin Zobel, and Tetsuya Sakai. Precision-at-ten considered redundant. In Sung-Hyon Myaeng, Douglas W. Oard, Fabrizio Sebastiani, Tat-Seng Chua, and Mun-Kew Leong, editors, *Proc. 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 695–696, Singapore, Singapore, July 2008c.

Alistair Moffat, William Webber, and Justin Zobel. Strategic system comparisons via targeted relevance judgments. In Charles L. A. Clarke, Norbert Fuhr, Noriko Kando, W. Kraaij, and A. de Vries, editors, *Proc. 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 375–382, Amsterdam, the Netherlands, July 2007b.

William Webber, Alistair Moffat, and Justin Zobel. Score standardization for robust comparison of retrieval systems. In MingFang Wu, Andrew Turpin, and Amanda Spink, editors, *Proc. 12th Australasian Document Computing Symposium*, pages 1–8, Melbourne, December 2007a.

Alistair Moffat, William Webber, and Justin Zobel. Load balancing for term-distributed parallel retrieval. In Susan Dumais, E. Efthimiadis, David Hawking, and Kalervo Järvelin, editors, *Proc. 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 348–355, Seattle, Washington, USA, August 2006.

William Webber and Alistair Moffat. In search of reliable retrieval experiments. In J. Kay, Andrew Turpin, and Ross Wilkinson, editors, *Proc. 10th Australasian Document Computing Symposium*, pages 26–33, Sydney, December 2005.

N. Lester, Alistair Moffat, William Webber, and Justin Zobel. Space-limited ranked query evaluation using adaptive pruning. In Anne H.H. Ngu, Masaru Kitsuregawa, Erich J. Neuhold, Jen-Yao Chung, and Quan Z. Sheng, editors, *Proc. 6th International Conference on Web Informations Systems Engineering*, volume 3806 of *Lecture Notes in Computer Science*, pages 470–477, New York, New York, USA, November 2005.

Peer-reviewed workshops

Justin Zobel, William Webber, Mark Sanderson, and Alistair Moffat. Principles for robust evaluation infrastructure. In *Proc. CIKM Data Infrastructures for Supporting Information Retrieval Evaluation*, pages 6:1–4, Glasgow, UK, October 2011.

William Webber. Re-examining the effectiveness of manual review. In *Proc. SIGIR Information Retrieval for E-Discovery Workshop*, pages 2:1–8, Beijing, China, July 2011.

William Webber, Alistair Moffat, and Justin Zobel. The effect of pooling and evaluation depth on metric stability. In Tetsuya Sakai, Mark Sanderson, and William Webber, editors, *Proc. 3rd International Workshop on Evaluating Information Access*, pages 7–15, Tokyo, Japan, June 2010c.

Timothy G. Armstrong, Justin Zobel, William Webber, and Alistair Moffat. Relative significance is insufficient: baselines matters too. In Shlomo Geva, Jaap Kamps, Carol Peters, Tetsuya Sakai, Andrew Trotman, and Ellen Voorhees, editors, *Proc. SIGIR 2009 Workshop on the Future of IR Evaluation*, pages 25–26, Boston, Massachusetts, USA, July 2009c.

Demonstrations

Timothy G. Armstrong, Alistair Moffat, William Webber, and Justin Zobel. EvaluatIR: An online tool for evaluating and comparing IR systems. In James Allan, Javed Aslam, Mark Sanderson, ChengXiang Zhai, and Justin Zobel, editors, *Proc. 32nd Annual International ACM SIGIR Conference on Research and Development in*

Information Retrieval, page 834, Boston, Massachusetts, USA, July 2009d.

Law and professional journals

William Webber and Jason R. Baron. Automation and sampling in e-discovery. *Internet Law Bulletin*, 16(7): 158–161, November 2013.

Non-peer-reviewed venues

William Webber, Falk Scholer, MingFang Wu, Xiuzhen Zhang, Douglas W. Oard, Phil Farrelly, Sandra Potter, Steven Dick, and Phill Bertolus. The Melbourne team at the TREC 2010 legal track. In Ellen Voorhees and Lori P. Buckland, editors, *Proc. 19th Text REtrieval Conference*, pages 49:1–12, Gaithersburg, Maryland, USA, November 2010d.

William Webber. Evaluating the effectiveness of keyword search. *IEEE Data Eng. Bull.*, 33(1):54–59, 2010a.

William Webber, V. N. Anh, and Alistair Moffat. The University of Melbourne in the Million Query Track of TREC 2007. In Ellen Voorhees and Lori P. Buckland, editors, *Proc. 16th Text REtrieval Conference*, page 86:1:5, Gaithersburg, Maryland, USA, November 2007b. NIST Special Publication 500-274.

V. N. Anh, William Webber, and Alistair Moffat. Melbourne University at the 2006 Terabyte Track. In Ellen Voorhees and Lori P. Buckland, editors, *Proc. 15th Text REtrieval Conference*, pages 103:1–8, Gaithersburg, Maryland, USA, November 2006. NIST Special Publication 500-272.

V. N. Anh, William Webber, and Alistair Moffat. Melbourne University 2005: Enterprise and Terabyte tracks. In Ellen Voorhees and Lori P. Buckland, editors, *Proc. 14th Text REtrieval Conference*, pages 114:1–8, Gaithersburg, Maryland, USA, November 2005. NIST Special Publication 500-266.

Yaniv Bernstein, Bodo Billerbeck, Steven Garcia, Nicholas Lester, Falk Scholer, Justin Zobel, and William Webber. RMIT university at TREC 2005: Terabyte and robust track. In Ellen Voorhees and Lori P. Buckland, editors, *Proc. 14th Text REtrieval Conference*, pages 78:1–12, Gaithersburg, Maryland, USA, November 2005. NIST Special Publication 500-266.

Bodo Billerbeck, Adam Cannane, Abhijit Chattaraj, Nicholas Lester, William Webber, Hugh Williams, John Yiannis, and Justin Zobel. RMIT university at TREC 2004. In Ellen Voorhees and Lori P. Buckland, editors, *Proc. 13th Text REtrieval Conference*, pages 55:1–8, Gaithersburg, Maryland, USA, November 2004. NIST Special Publication 500-261.

Theses

William Webber. *Measurement in Information Retrieval Evaluation*. PhD thesis, University of Melbourne, Melbourne, Australia, September 2010b.

William Webber. Design and implementation of a pipelined distributed information retrieval architecture. Master's thesis, University of Melbourne, Melbourne, Australia, June 2007.

PROFESSIONAL ACTIVITIES

Journal editorships

Information Retrieval
Member of Editorial Board

2011 to 2014

Conference and workshop chairs

SIGIR area co-chair for evaluation (2013); EVIA co-chair (2010–2013); AIRS area co-chair for evaluation (2010).

Conference and workshop program committees

IJCNLP (2013); IJCAI (2011); SIGIR (2010–2012); CIKM (2009, 2011); AIRS (2012); ADCS (2007, 2009–2010).

Journal reviewing

ACMTOIS; JASIST; IEEE TKDE; TACL; The Computer Journal; Information Processing Letters.

TALKS

E-Discovery Protocols for Predictive Coding August, 2013

Talk

Presented iterative e-discovery process with termination cost guarantees to the Victorian Society for Computers and the Law.

Best Practices for Leveraging Predictive January, 2013

Technologies

Webinar

Spoke on evaluation and certification best practice in e-discovery. Co-panelist with Ralph Losey, Esq. (Partner and National E-Discovery Counsel, Jackson Lewis LLC.)

SIGIR 2012 Panel on the Use of Proprietary Data in July, 2012

Information Retrieval Research

Panelist

Spoke in panel discussion on whether and how proprietary data should be used in public research work. Other panelists were Ricardo Baeza-Yates (Yahoo! Research), Susan Dumais (Microsoft Research), Norbert Fuhr (University of Duisburg-Essen), William Hersh (Origin Health and Science University), and Maarten de Rijke (University of Amsterdam).

PROFESSIONAL MEMBERSHIPS

Association for Computing Machinery (ACM)

- ACM Special Interest Group on Information Retrieval (SIGIR).

Statistical Society of Australia Inc. (SSAI).

Australian Computer Society (ACS).

American Society of Information Science and Technology (ASIST).

American Statistical Association (ASA).

Victorian Society for Computers and the Law (VSCL).

Last updated: 15th March, 2014